



University of Pennsylvania
ScholarlyCommons

Operations, Information and Decisions Papers

Wharton Faculty Research

7-2010

Fostering IPv6 Migration Through Network Quality Differentials

Guerin Roch

kartik Hosanagar
University of Pennsylvania

Follow this and additional works at: https://repository.upenn.edu/oid_papers

 Part of the [Other Economics Commons](#)

Recommended Citation

Roch, G., & Hosanagar, k. (2010). Fostering IPv6 Migration Through Network Quality Differentials. *ACM SIGCOMM Computer Communication Review*, 40 (3), 17-25. <http://dx.doi.org/10.1145/1823844.1823847>

This paper is posted at ScholarlyCommons. https://repository.upenn.edu/oid_papers/244
For more information, please contact repository@pobox.upenn.edu.

Fostering IPv6 Migration Through Network Quality Differentials

Abstract

Although IPv6 has been the next generation Internet protocol for nearly 15 years, new evidences indicate that transitioning from IPv4 to IPv6 is about to become a more pressing issue. This paper attempts to quantify if and how such a transition may unfold. The focus is on "connectivity quality," e.g., as measured by users' experience when accessing content, as a possible incentive (or disincentive) for migrating to IPv6, and on "translation costs" (between IPv6 and IPv4) that Internet Service Providers will incur during this transition. The paper develops a simple model that captures some of the underlying interactions, and highlights the ambiguous role of translation gateways that can either help or discourage IPv6 adoption. The paper is an initial foray in the complex and often puzzling issue of migrating the current Internet to a new version with which it is incompatible.

Keywords

IPv6, migration, incentives, quality

Disciplines

Other Economics

Fostering IPv6 Migration Through Network Quality Differentials*

Roch Guérin
Elec. & Sys. Eng, U. Pennsylvania
guerin@ee.upenn.edu

Kartik Hosanagar
Wharton, U. Pennsylvania
kartikh@wharton.upenn.edu

ABSTRACT

This paper develops a simple model that explores the extent to which differences in connectivity quality can be effective incentives (or disincentives) towards migrating to an IPv6 Internet. The focus is on “translation” costs that service providers will incur if slow adoption of IPv6 persists in the current IPv4 Internet, while continued growth in the Internet’s user base increasingly forces the use of IPv6 addresses for new users. The model elucidates how differences in connectivity quality, *i.e.*, between native IPv6, native IPv4, and what is achievable through translation devices or IPv6↔IPv4 gateways, affect IPv6 adoption in the current Internet. It highlights how improving native IPv6 connectivity quality, and obviously publicizing that better quality, may provide sufficient incentives to foster a more rapid migration to IPv6. It also reveals the often ambiguous role of gateways that can help or discourage IPv6 adoption. The paper represents an initial foray in the complex and often frustrating problem of migrating the current Internet to a next generation technology with which it is incompatible.

Categories and Subject Descriptors

C.2.3 [Network Operation]: Public networks

C.2.6 [Internetworking]: Standards

General Terms

Economics, Management, Standardization

Keywords

IPv6, migration, incentives, quality

1. BACKGROUND AND MOTIVATIONS

IPv6, the next generation Internet Protocol, was standardized about fifteen years ago [3] to address a number of known deficiencies in the current (IPv4) version of the protocol. In particular, it increased address size from 32 bits to 128 bits to ensure that as the Internet popularity and the number of devices connected to it grew, it would not run out of addresses¹ and in the process curtail growth. In spite of this early standardization and the impending exhaustion of the IPv4 address pool (various estimators, *e.g.*, <http://www.potaroo.net/tools/ipv4>, put the exhaustion of the address pool of the Internet Assigned Numbers

*This work was supported by NSF grant CNS-0721610.

¹A total of $2^{32} \approx 4.3$ billions addresses may seem plentiful, but we are fast approaching this limit.

Authority -IANA- to September 2011, and the exhaustion of the Regional Internet Registry -RIR- pools to about a year later), IPv6 adoption has been mostly marginal (see [8] or <http://mnlab-ipv6.seas.upenn.edu/monitor/index.html> for how much Internet content is accessible over IPv6).

There are many explanations for this lack of progress (see [7] for an insightful discussion), including few real incentives for IPv6 deployment (it does not really offer any substantial benefits over IPv4 except for its larger address space), and until recently limited support for IPv6 in many end-systems². Irrespective of the reasons, the slow deployment of IPv6 means that it now faces a formidable incumbent, the current IPv4 Internet, with which it is *incompatible*³. This in itself represents a significant hurdle for an eventual migration to an IPv6 Internet. Specifically, even if we ultimately proceed with assigning IPv6 addresses to new devices (bridge gap options such as private addresses have, among others, translation costs that keep rising with usage), this only accounts for half the issue. In particular, it does not address the challenge of convincing *existing* IPv4 devices to also adopt IPv6, *i.e.*, acquire an IPv6 address in addition to their IPv4 address.

Given the size of the Internet, making the installed IPv4 base, and in particular *content*, accessible through IPv6 is critical on multiple accounts. First and foremost, access to content and other users is what defines the value of the Internet, and because IPv4 and IPv6 are incompatible, an IPv6-only device, *i.e.*, a device without an IPv4 address, cannot directly access any device or content that is only reachable over IPv4. Specifically, an IPv6-only device seeking to connect to a “site” would as usual query the Domain Name System (DNS) for the IP address associated with the site’s name, but would do so requesting an address of type AAAA (quad-A). In the absence of a registered IPv6 address for the site, DNS would inform the user that the site is unreachable (over IPv6). IPv6 Internet connectivity is, therefore, of little or no value to (IPv6) users without gateways or “translation” devices, *e.g.*, [5], that provide access to the IPv4 Internet. In other words, the lack of IPv6 adoption in the current IPv4 Internet mandates the deployment of gateways to entice new users to accept IPv6. Furthermore, limited IPv6 adoption in the IPv4 Internet also increases the capacity requirements and, therefore, costs of those gateways. This is because the volume of traffic traversing IPv6↔IPv4

²Native IPv6 support only became truly available in the Windows operating system with the Vista release.

³In hindsight, this is probably the single biggest mistake made during the specification of IPv6.

gateways is “proportional” to both the number of IPv6-only users and the number of IPv4-only (not accessible over IPv6) sites those users wish to access.

As a result, fostering the adoption of IPv6 in the current IPv4 Internet is vital to an eventual migration to an IPv6 Internet; something that many see as desirable (see [6] for a discussion on various alternatives). This is, however, not without challenges as the many failed efforts to promote or even mandate⁴ IPv6 have shown. In particular, IPv6 adoption depends in part on providing the right incentives to the different stake-holders, *i.e.*, Internet users, Internet service providers (ISPs), and Internet content providers (ICPs). ISPs play an especially critical role as owners of the infrastructure connecting Internet users to ICPs. Through pricing of Internet services, ISPs can influence the adoption of IPv6 by both new and existing (IPv4) Internet users, and conversely through tuning the quality differential of the connectivity they offer over IPv6 and IPv4 they may be able to influence the adoption of IPv6 by ICPs. For example, Google has made itself accessible over IPv6, but through a different “service”, *i.e.*, as <http://ipv6.google.com> and not the standard <http://www.google.com>, *except* for users connected to an ISP certified by Google as having *good* IPv6 connectivity (see <http://www.google.com/intl/en/ipv6/>).

Our goal in this paper is not to provide a comprehensive investigation of all the possible incentives available for fostering migration to IPv6. Instead, we focus on how differences in *connectivity quality* can influence IPv6 adoption and the volume of translation traffic that ISPs should provision for. For that purpose, we develop and analyze a minimalist model that highlights a number of interesting phenomena that arise as a function of differences in connectivity quality. We find that gateways can often play an ambiguous role and can help or discourage IPv6 adoption. Further, we highlight how improving native IPv6 connectivity quality, and obviously publicizing that better quality, may provide sufficient incentives to foster a more rapid migration to IPv6.

The rest of the paper is structured as follows. Section 2 introduces our minimalist model, its notations and the assumptions on which it relies. The model is analyzed and discussed in Section 3, which in particular demonstrates how translation traffic can be kept bounded by varying the relative quality of the different connectivity options. Section 4 summarizes key findings of the model and their implications for IPv6 adoption, and points to possible extensions to generalize the model and/or eliminate some of the simplifying assumptions on which it relies.

2. MODEL AND NOTATION

As alluded to earlier, our model involves three “players”: users, a service provider (ISP) and content providers (ICPs). Its goals are to assess how connectivity quality choices made by the ISP affect IPv6 adoption by ICPs, and consequently the volume of translation traffic the ISP needs to provision for. Users access ICPs that derive revenue from users, so that ICPs make IPv6 adoption decisions based on how the connectivity quality choices of the ISP affect user access. Subsection 2.1 details the roles and interactions of the three players, while Subsection 2.2 articulates the resulting model.

⁴June 30, 2008 was the deadline set by the US Government’s Office of Management and Budget for federal agencies to be running IPv6 [2], and there have been numerous similar initiatives in other countries as reported in [9].

2.1 Assumptions and Notation

Consider an ISP (more generally a set of ISPs) that because of shortage of IPv4 addresses has started allocating IPv6 addresses to new users. Implicit in this choice is a preference for IPv6 over private IPv4 addresses. As mentioned earlier, this is because private addresses incur similar “translation costs” as IPv6 addresses when connecting to the public IPv4 Internet, *i.e.*, from a private IPv4 address to a public one and back, without the possibility of an eventual migration to a translation-free environment. In other words, while private addresses may offer short-term benefits, *i.e.*, the use of a familiar technology, their long-term costs keep growing with the size of the Internet. In contrast, IPv6 incurs short-term deployment and training costs, but has the potential for much lower long-term costs. Hence, it can be argued to be the better option if its long-term benefits can be realized. Investigating how this can be accomplished is one of the paper’s motivations. Obviously, in this model the ISP is also assumed to operate a network that is both IPv4 and IPv6 capable.

The ISP has two categories of users: Existing users that have been allocated an IPv4 address (and possibly also an IPv6 address), and new users that only have an IPv6 address. Users have no control over the type of IP address they receive and are not decision makers in the model. In particular, we ignore possible competition for new users among ISPs, at least to the extent that all ISPs are assumed to have run out of IPv4 addresses and to have started allocating IPv6 addresses to all new users. Alternatively, one may assume that most users would not have the technical wherewithal to differentiate between an IPv4 or an IPv6 address, especially if this is “hidden” through the deployment of an ISP owned access devices.

The main effect of users is their ability to access Internet content, which as we shall see can drive the IPv6 adoption decisions of ICPs as well as the connectivity quality options that the ISP selects. The size of the IPv4 and IPv6 user populations are denoted as x_4 and $x_6(t)$, respectively, where x_4 is fixed (the user population when the ISP ran out of IPv4 addresses), and $x_6(t)$ is a non-decreasing function of time (t) with $x_6(0) = 0$. We further assume that $x_6(t)$ is an exogenous function so that the growth in the ISP’s user base is independent of the connectivity quality choices it makes. In other words and consistent with the above discussion, while differences in the connectivity quality of IPv6 users may exist, they are not significant enough to affect user adoption decisions. In addition to new users that are assigned only an IPv6 address, a fraction $\alpha(t)$, $0 \leq \alpha(t) \leq 1$, of existing x_4 IPv4 users is assumed to be also IPv6 accessible, *i.e.*, have been assigned an IPv6 address by the ISP in addition to their IPv4 address. $\alpha(t)$ is a decision variable of the ISP that controls when to also allocate an IPv6 address to existing IPv4 users (IPv6 addresses are not scarce and many user end-systems nowadays support both IPv4 and IPv6).

Users are assumed homogeneous in how they access *content providers* (ICPs), and in the amount of traffic they generate. ICPs are also homogeneous in how users access them and in the amount of user traffic they sink and source. This is obviously an over-simplification as ICPs vary in popularity, *e.g.*, [1], which influences traffic volume, but should not significantly affect our main conclusions, *e.g.*, IPv6 adoption by a popular site is essentially equivalent to adoption by multiple sites in our simple model. Transition points may

shift, but general behaviors should remain similar. All ICPs own an IPv4 address (registered with DNS), and can also register an IPv6 address to enable IPv6 connectivity. Note that implicit here is the assumption that all ICPs are IPv6 reachable through the ISP and more generally the Internet⁵. For simplicity, the model ignore IPv6-only ICPs (their number is anyhow marginal). IPv6 users that seek to connect to IPv4-only ICPs must rely on translation devices such as [5] (the ISP is responsible for this translation function), while they can connect natively to ICPs that have registered an IPv6 address. The fraction of ICPs that have decided to make themselves accessible over both IPv4 and IPv6 is denoted as $\beta(t)$, $0 \leq \beta(t) \leq 1$. As discussed below, the variable $\beta(t)$ is endogenous and driven by the IPv6 adoption process of ICPs.

There are three possible types of network connectivity and associated quality levels: q_{44} , $q_{64}(t)$ and $q_{66}(t)$. They correspond to the three possible combinations of user and ICP reachability, namely, IPv4 \leftrightarrow IPv4, IPv6 \leftrightarrow IPv4, IPv6 \leftrightarrow IPv6, respectively. Connectivity quality of the IPv4 Internet is assumed fixed, *i.e.*, $q_{44} = 1$, while connectivity quality through translation devices, $q_{64}(t)$, and native IPv6 connectivity, $q_{66}(t)$, can vary with time, *e.g.*, $q_{64}(t)$ can decrease as translation boxes become more heavily loaded and conversely $q_{66}(t)$ can improve as technology and skills mature. We assume that $q_{64}(t) \leq q_{44}$, *i.e.*, translation can only lower the quality of accessing content over the IPv4 Internet. However, there is no such constraint on the quality of IPv6 connectivity, $q_{66}(t)$, that can be better or worse than IPv4 connectivity with or without translation. For example, depending on the capabilities of network devices, better IPv6 quality could be realized in a number of different ways, *e.g.*, by giving precedence to IPv6 packets, or allocating a larger bandwidth share to IPv6 traffic on router interfaces, etc. The quantities $q_{64}(t)$ and $q_{66}(t)$ are decision variables of the ISP that is, therefore, assumed to “control” the quality of the connectivity between users and content providers. This is clearly a simplifying assumption in the model, as connectivity quality is an end-to-end property that often involves resources (other ISPs) beyond the control of our target ISP (or set of ISPs). Our notion of connectivity quality should, therefore, be viewed as an “average” measure that may also incorporate limits on how much an individual ISP (or set of ISPs) can affect its value.

The network provider (ISP) incurs costs that arguably depend on both the quality of the connectivity it offers and the volume of traffic it carries. More importantly in the context of this study, it incurs translation costs to allow new IPv6-only users to connect to content providers (ICPs) that are accessible only over IPv4. These costs increase with the volume of translation (IPv6 \leftrightarrow IPv4) traffic that needs to be handled, and are a focus of this paper. In this study, we ignore connectivity costs or assume that they remain approximately constant across connectivity options and instead focus on translation costs. This is clearly an approximation, but one that is reasonable given our focus on exploring the

impact of differences in connectivity quality on IPv6 adoption. Nevertheless, as discussed in Section 4, extending the model to incorporate a joint optimization of connectivity quality and translation costs is clearly of interest and a topic we plan to further investigate.

ICP revenues grow with the number of (IPv4 and IPv6) users in the network and their “activity” (the traffic they generate). The network “activity” of users is assumed to be proportional to the quality q of their network connectivity. Correspondingly, for a given user population of size x , an ICP’s revenue is of the general form $R \sim f(qx)$. As with the volume of traffic they sink and source, ICPs are assumed homogeneous in the revenue they derive from users, and for simplicity we assume $R \sim qx$. These revenues are offset by ICP connectivity costs. The cost component of interest in the context of this paper is the configuration cost for making themselves accessible to users. This cost depends on whether they are accessible over IPv4 only, or over both IPv4 and IPv6, *i.e.*, is of the form $c_4^{(i)}$ or $c_4^{(i)} + c_6^{(i)}$, respectively, for ICP i . To capture the disparity that exists among content providers in both awareness about IPv6 and technical know-how on how to deal with it, the component $c_6^{(i)}$ is heterogeneous across ICPs. This heterogeneity is captured through a parameter θ_i , which is taken to be uniformly distributed between 0 and 1. For simplicity, IPv4 configuration costs are taken to be constant across ICPs (IPv4 is a mature technology) and set equal to 0. As a result the configuration cost of ICP i is $c_4^{(i)} + c_6^{(i)} = 0 + \theta_i c_6$. ICPs will, therefore, decide on a connectivity option (IPv4 only or both IPv4 and IPv6) as a function of its cost and how it affects their revenue. This decision process is formalized in the next section.

2.2 Model Formulation

2.2.1 ISP Decision Process and Variables

As stated in the previous section, ISPs’ costs include both connectivity and translation costs. Because of our focus on migration issues, we ignore the former and concentrate on the latter that are assumed proportional to the volume of translation traffic, which at time t is of the form

$$T(t) \sim x_6(t) (1 - \beta(t)) \quad (1)$$

Eq. (1) reflects the dependency of $T(t)$ on the level of IPv6 adoption by ICPs⁶. The ISP’s decision process involves choosing connectivity quality options, $q_{64}(t)$ and $q_{66}(t)$, as well as $\alpha(t)$ (the extent to which existing IPv4 users are also provided with IPv6 addresses), to maximize profit or minimize the amount of translation traffic it needs to handle (maximize $\beta(t)$). Alternatively, as discussed in Section 3.1.2, the ISP may simply wish to keep translation traffic volume below a certain value, *e.g.*, so as not to exceed the capacity of the translation devices it has deployed. In either case, the ISP’s decisions will be based on how they affect ICPs IPv6 adoption, which we describe next.

2.2.2 ICP Decision Process and Variables

ICPs derive revenues from users accessing them and incur costs when they decide to become IPv6 accessible. ICPs, therefore, evaluate connectivity options and select the one

⁶If all content was accessible over IPv6 ($\beta(t) = 1$) there would be no need for translation; the original assumption when IPv6 was first standardized.

⁵Although a large number of ASes are still not using IPv6, many of the largest ISPs are IPv6 enabled, *e.g.*, see <http://bgp.he.net/ipv6-progress-report.cgi>, and connectivity to IPv6 enabled sites is broadly available (in over 6.5 million tests, only 0.13% identified problems for users trying to reach a web site after it became IPv6 reachable, see <http://ipv6test.max.nl>).

with the highest profit (revenue minus cost). We consider two different models for ICPs' profits. Model 1 assumes that IPv6 addresses have precedence over IPv4 addresses, *i.e.*, if content is accessible over both IPv4 and IPv6, users that have the option to select either, *i.e.*, the $\alpha(t)x_4$ users configured with both IPv4 and IPv6 addresses, will select IPv6 connectivity. This is consistent with the default policy of [4]. Model 2 considers an alternative setting where users with both IPv4 and IPv6 connectivity select the one that affords the better quality. This corresponds to a rational decision process by users or providers of end-user devices⁷, based on awareness of quality differentials across connectivity options. For each model, expressions for ICP i 's profits under different connectivity options are as follows

Model 1: IPv6 has precedence:

$$\Pi_4^{(i)} = x_4 + q_{64}(t)x_6(t) \quad (2)$$

$$\begin{aligned} \Pi_{46}^{(i)} &= x_4(1 - \alpha(t)) \\ &+ (\alpha(t)x_4 + x_6(t))q_{66}(t) - \theta_i c_6 \end{aligned} \quad (3)$$

Eq. (2) is the ICP profit if it remains only IPv4 accessible, while Eq. (3) gives its profit once it also becomes IPv6 accessible. Eq. (3) reflects the precedence of IPv6 over IPv4.

Model 2: Quality has precedence:

$$\Pi_4^{(i)} = x_4 + q_{64}(t)x_6(t) \quad (4)$$

$$\Pi_{46}^{(i)} = x_4 + x_6(t)q_{66}(t) - \theta_i c_6 \quad \text{if } q_{66}(t) \leq 1 \quad (5)$$

$$\begin{aligned} \Pi_{46}^{(i)} &= x_4(1 - \alpha(t)) + (\alpha(t)x_4 + x_6(t))q_{66}(t) \\ &- \theta_i c_6 \quad \text{if } q_{66}(t) > 1 \end{aligned} \quad (6)$$

Eq. (4) is identical to Eq. (2). Eq. (5) gives the ICP profit once it becomes IPv6 accessible with IPv6 quality lower than that of IPv4, while Eq. (6) considers the case where IPv6 quality is higher than that of IPv4. In both models, ICP i decides to become IPv6 accessible only if it yields a higher profit, *i.e.*, $\Pi_{46}^{(i)} > \Pi_4^{(i)}$.

Let $\delta_{46}^{(i)}(t)$ denote the variable that reflects the decision of ICP i to become IPv6 accessible or not, *i.e.*, $\delta_{46}^{(i)}(t) = 1$ if it does and 0 otherwise. Then under both models, we have

$$\delta_{46}^{(i)}(t) = \begin{cases} 1 & , \text{ if } \Pi_4^{(i)} < \Pi_{46}^{(i)} \\ 0 & , \text{ otherwise} \end{cases} \quad (7)$$

3. ANALYSIS AND DISCUSSION

3.1 Model 1 – IPv6 has Precedence

3.1.1 ICP IPv6 Adoption

From Eqs. (2), (3), and (7), together with the assumption that θ_i is uniformly distributed in $[0, 1]$, we get the following expression for the level $\beta(t)$ of IPv6 adoption by ICPs

$$\beta(t) = \left(\frac{\alpha(t)x_4(q_{66}(t) - 1) + x_6(t)(q_{66}(t) - q_{64}(t))}{c_6} \right)_{[0,1]}, \quad (8)$$

⁷If IPv6 quality is well-known to be poor, providers of operating systems may ship them configured to always prefer IPv4 whenever there is a choice, or conversely system administrators may choose this as the default configuration for systems they manage.

where we have used the notation $(x)_{[0,1]}$ to indicate the projection of x on the interval $[0, 1]$.

Eq. (8) provides intuitive confirmation of the impact of different parameters on IPv6 adoption by ICPs.

The first intuitive finding from Eq. (8), is the fact that without quality differentials IPv6 adoption may never take off, *i.e.*, $\forall t \beta(t) = 0$ if $q_{64}(t) = q_{66}(t) = q_{44} = 1$.

Another finding is that IPv6 adoption grows the better the quality of native IPv6 connectivity (q_{66}), but remains at 0 unless IPv6 affords better performance than what is achievable from going through translation devices, *i.e.*, $\beta = 0$ if $q_{66} \leq q_{64} \leq 1$. That is, no ICP will adopt IPv6 if translation quality is better than native IPv6 quality. This represents somewhat of a dilemma for ISPs, which may struggle with their early IPv6 deployment, *i.e.*, $q_{66}(t)$ is likely to initially be less than $q_{44} = 1$, while providing reasonable translation quality, *i.e.*, $q_{64}(t) \approx q_{44} = 1$, may at first be relatively easy as translation traffic volume $(1 - \beta(t))x_6(t)$ will be low. There will, therefore, be no adoption of IPv6 by ICPs until IPv6 connectivity quality surpasses that of translation. At this point, some ICPs become incentivized to adopt IPv6, *i.e.*, $\Pi_{46}^{(i)} > \Pi_4^{(i)}$ for some i .

Additionally, if IPv6 connectivity is no better than IPv4 connectivity, *i.e.*, $q_{66} \leq q_{44} = 1$, it is harmful to ICP IPv6 adoption to start providing existing IPv4 users with IPv6 addresses, *i.e.*, increasing $\alpha(t)$ results in lower IPv6 adoption (smaller β values) by ICPs. This is due to Model 1's assumption that users with both IPv4 and IPv6 addresses give precedence to IPv6 access when that option is available. This forces users with IPv4 and IPv6 addresses to connect to IPv6 accessible sites using IPv6, and in the process experience lower quality than if they had used IPv4. This represents a disincentive for ICPs to become IPv6 accessible, in spite of the fact that IPv6 offers IPv6-only users better connectivity quality than what they experience with translation devices. Specifically, no ICP chooses to become IPv6 accessible, *i.e.*, $\beta(t) = 0$, if

$$\alpha(t) \geq \frac{x_6(t)(q_{66}(t) - q_{64}(t))}{x_4(1 - q_{66}(t))} \quad \text{and } q_{66}(t) < 1$$

As a result, the ISP's best strategy to maximize IPv6 adoption is to avoid enabling IPv6 access for existing IPv4 users, *i.e.*, keep $\alpha(t) = 0$, if $q_{66} \leq 1$. Conversely, when IPv6 connectivity becomes better than IPv4 connectivity, making existing IPv4 users accessible over IPv6 becomes beneficial to ICP IPv6 adoption, which grows with $\alpha(t)$. In general, IPv6 adoption by ICPs, $\beta(t)$, is of the form

$$\beta(t) = \begin{cases} 0, & \text{if } q_{66}(t) < q_{64}(t) \leq 1 & (a) \\ \left(\frac{\alpha(t)x_4(q_{66}(t) - 1) + x_6(t)(q_{66}(t) - q_{64}(t))}{c_6} \right)_{[0,1]}, & \text{otherwise} & (b) \end{cases} \quad (9)$$

The main (intuitive) findings from the above discussion are summarized in Proposition 1.

Proposition 1 *IPv6 adoption by Internet Content Providers (ICPs) is influenced as follows by differences in the levels of IPv4 and IPv6 connectivity quality*

- ICPs will not be incented to become IPv6 accessible as long as IPv6 connectivity is no better than what translation devices offer, *i.e.*, $\beta(t) = 0$, if $q_{66}(t) \leq q_{64}(t)$;

- IPv6 adoption by ICPs will grow with the number of IPv6-only users as soon as IPv6 connectivity quality exceeds that of translation devices, i.e., $\frac{\partial \beta}{\partial \alpha} > 0$ if $q_{66}(t) > q_{64}(t)$;
- As long as IPv6 connectivity is not equal to or better than IPv4 connectivity, it is best not to allow existing IPv4 users the option of accessing content over IPv6, i.e., $\frac{\partial \beta}{\partial \alpha} < 0$ if $q_{66}(t) < 1$;
- Conversely, once IPv6 quality exceeds that of the IPv4 Internet, IPv6 adoption by ICPs will benefit from making existing IPv4 users IPv6 enabled, i.e., $\frac{\partial \beta}{\partial \alpha} > 0$ if $q_{66}(t) > 1$.

3.1.2 Managing Translation Traffic Growth

Another aspect of interest is the volume of translation traffic that an ISP needs to deal with, and therefore provision for. Using Eqs. (1) and (9), we can investigate the effect of different strategies on this quantity.

Consider a scenario in which the ISP has deployed translation devices that have an aggregate capacity of a . The ISP is then interested in exploring means to keep translation traffic below that value even as the number of new IPv6 users, $x_6(t)$, increases. In other words, can the ISP “control” the volume of translation traffic through its three decision variables, $\alpha(t)$, $q_{64}(t)$, and $q_{66}(t)$, knowing how they affect ICPs’ decisions as predicted by Eq. (9). Note that this control is indirect, *i.e.*, through its influence on ICP decisions, and is not based on throttling or rate-limiting the amount of translation traffic that IPv6 users originate. This is control at the time-scale of provisioning, when the ISP determines if and by how much to upgrade its translation capacity.

The volume of translation traffic can be obtained from combining Eq. (1) and Eq. (9). Following Eqs. (9)(a) and (b), we consider three distinct configurations.

I. $q_{66}(t) < q_{64}(t) \leq 1$.

From Eq. (9) this results in $\beta(t) = 0$, *i.e.*, no ICPs make themselves IPv6 accessible. Translation traffic volume is, therefore, directly proportional to the number of new IPv6 users, $x_6(t)$, and the number of existing IPv4 users to which the ISP has also provided an IPv6 address, $\alpha(t)x_4$. As a result, if the ISP is unable to provide IPv6 connectivity quality that is better than what is achievable through translation, its only option to keep the volume of translation traffic below a is simply to ensure that the total (to all ICPs) combined traffic from new IPv6 users ($x_6(t)$) and existing IPv4 users that have been provided with an IPv6 address ($\alpha(t)x_4$) does not exceed a . Obviously, this is best achieved by not configuring IPv6 addresses on any existing IPv4 users, *i.e.*, by setting $\alpha(t) = 0$, and remains feasible as long as there are not too many new IPv6 users ($x_6(t) \leq a$). As soon as the traffic from new IPv6 users exceeds a , *i.e.*, $x_6(t) > a$, the ISP has no choice but to improve IPv6 quality to keep translation traffic volume below a .

In other words, in **Configuration I**, the ISP’s best strategy is to set $\alpha(t) = 0$, but the strategy is viable, *i.e.*, ensures $T(t) \leq a$, only as long $x_6(t) \leq a$.

The next configuration considers the case where the ISP has improved IPv6 connectivity quality beyond that achievable through translation, but not yet so that it surpasses IPv4 quality, *i.e.*, $q_{64}(t) \leq q_{66}(t) < q_{44} = 1$.

II. $q_{64}(t) \leq q_{66}(t) < 1$.

As in the previous scenario, the ISP’s best strategy to minimize translation traffic volume is to set $\alpha(t) = 0$. As

discussed in the context of Eq. (8), this is because as long as the quantity $(q_{66}(t) - 1)$ is negative, allowing existing IPv4 users native IPv6 access has a negative effect on $\beta(t)$. Increasing $\alpha(t)$, therefore, contributes additional translation traffic.

There are, however, additional conditions that need to be satisfied to keep translation traffic volume below a when $x_6(t) > a$. In particular, IPv6 connectivity, $q_{66}(t)$, needs to exceed translation quality, $q_{64}(t)$, by a certain margin. We explore this issue next.

From Eq. (9) and assuming $q_{64}(t) \leq q_{66}(t) < 1$ (and $\alpha(t) = 0$), ensuring $T(t) \leq a$ yields the following relation between $q_{66}(t)$ and $q_{64}(t)$:

$$q_{66}(t) \geq q_{64}(t) + \frac{c_6}{x_6(t)} \left(1 - \frac{a}{x_6(t)} \right). \quad (10)$$

For a given value of $q_{64}(t)$, the right-hand-side of Eq. (10) is easily found to have a maximum for $x_6(t) = 2a$, which based on Eq. (1) is also where translation traffic peaks (at a value of a). Hence, to keep translation traffic below a for all values of $x_6(t)$, IPv6 quality must be better than a minimum value $q_{66}^{\min}(t)$ that satisfies.

$$q_{66}^{\min}(t) = q_{64}(t) + \frac{c_6}{4a}. \quad (11)$$

As expected, $q_{66}^{\min}(t)$ grows with c_6 , the cost of IPv6 configuration for ICPs, and decreases with the provisioned capacity of translation devices a . When combined with Eq. (9) and keeping $\alpha(t) = 0$, Eq. (11) also tells us that if $q_{66}(t) \geq q_{66}^{\min}(t)$, all ICPs will have become IPv6 accessible (the IPv4 Internet will have fully migrated to IPv6) once $x_6(t) = 4a$.

From Eq. (11), we see that for the ISP to be able to maintain $T(t) \leq a$ as the IPv6 user base grows without making IPv6 quality better than that of IPv4, *i.e.*, remain in **Configuration II**, translation quality must satisfy

$$q_{64} < 1 - \frac{c_6}{4a}. \quad (12)$$

The low translation quality is then sufficient to entice more ICPs to become IPv6 accessible as the number of IPv6 users increases. As before, it is best to set $\alpha(t) = 0$ to avoid disincentives associated with worsening connectivity quality for IPv4 users that also have an IPv6 address, *i.e.*, because $q_{66}(t) < 1$. How low translation quality needs to be depends on the ratio $\frac{c_6}{4a}$, and too low a translation quality may negatively affect the demand for IPv6, *i.e.*, $x_6(t)$; a factor that the current model does not account for.

It may, therefore, not always be possible to keep translation quality low enough, *i.e.*, we may need to have $q_{64} \geq 1 - \frac{c_6}{4a}$. Exploring this scenario is what we consider next. In particular, because keeping translation traffic below a will then at some point require improving IPv6 quality beyond that of IPv4 (or alternatively upgrade translation capacity to $a' > a$ so that Eq. (12) is now satisfied), it is of interest to determine when this happens, *i.e.*, for what value of $x_6(t)$.

Combining the condition $q_{66}(t) < 1$ with Eq. (10) yields the following inequality

$$(1 - q_{64})x_6^2 - c_6x_6 + ac_6 > 0. \quad (13)$$

When $q_{64} \geq 1 - \frac{c_6}{4a}$, Eq. (13) has two real-valued roots of

the form

$$x_6^{(1)} = \frac{c_6 - \sqrt{c_6^2 - 4ac_6(1 - q_{64})}}{2(1 - q_{64})} \quad (14)$$

$$x_6^{(2)} = \frac{c_6 + \sqrt{c_6^2 - 4ac_6(1 - q_{64})}}{2(1 - q_{64})} \quad (15)$$

As q_{64} varies in the range $(1 - \frac{c_6}{4a}, 1)$, the first root varies from $2a^-$ down to a^+ , while the second root varies from $2a^+$ to ∞ . Furthermore, $f(x)$ is negative in the range $(x_6^{(1)}, x_6^{(2)})$ and positive outside. This implies that Eq. (10) is satisfied only in the range $x_6(t) \in [a, x_6^{(1)})$. When $x_6(t)$ exceeds $x_6^{(1)}$, it becomes necessary for IPv6 quality to exceed that of IPv4, *i.e.*, $q_{66}(t) \geq 1$, to ensure that the volume of translation traffic remains below a . In other words, better translation quality ($q_{64} \geq 1 - \frac{c_6}{4a}$) delays IPv6 adoption by ICPs to the extent that IPv6 quality needs to improve beyond that of IPv4 to keep the translation traffic below a . This transition occurs when the number of IPv6-only users, $x_6(t)$, reaches $x_6^{(1)}$, with that value decreasing from $2a$ to a as translation quality approaches IPv4 quality. Consequently, full IPv6 adoption by ICPs ($\beta(t) = 1$) will also not happen until after that transition has occurred. Investigating when this happens is addressed in **Configuration III**, where we explore the impact of improved IPv6 connectivity on both translation traffic volume and full IPv6 adoption.

III. $q_{66}(t) \geq 1$.

From Eq. (1) and Eq. (9)(b), we obtain the following condition to ensure $T(t) \leq a$

$$(q_{66} - q_{64})x_6^2 - (c_6 - \alpha x_4(q_{66} - 1))x_6 + ac_6 \geq 0, \quad (16)$$

where for ease of notation we omitted dependency on t .

Because we are in the regime $q_{66}(t) \geq 1$, *i.e.*, IPv6 quality has been improved beyond that of the current IPv4 Internet, we want to identify conditions under which this is sufficient to always keep translation traffic volume below a . Eq. (16) implies that this requires

$$q_{66}(t) \geq \frac{q_{64}(t)x_6^2(t) + (c_6 + \alpha(t)x_4)x_6(t) - ac_6}{x_6^2(t) + \alpha(t)x_4x_6(t)} \quad (17)$$

Note that when $\alpha(t) = 0$, Eq. (17) simplifies to

$$q_{66}(t) \geq q_{64}(t) + \frac{c_6}{4a} \quad (18)$$

This is consistent with Eq. (11), and restates that when no existing IPv4 users are capable of accessing IPv6 content natively, ensuring that IPv6 connectivity quality exceeds the threshold of Eq. (18) is sufficient to ensure that translation traffic always remains below a . As expected, this threshold is higher when c_6 is high (high migration costs for ICPs) and a is low (limited translation capacity).

Conversely, we easily see from Eq. (17) that increasing $\alpha(t)$ allows the inequality to be met with a smaller value of $q_{66}(t)$ for all $x_6(t)$. In other words, when IPv6 connectivity quality is better than that of IPv4, making more IPv4 users IPv6 capable is beneficial to keeping the volume of translation traffic low. More generally, from Eq. (17) we can compute the value of $x_6^*(t)$ for which $q_{66}(t)$ realizes its maximum value, namely

$$x_6^*(t) = \frac{ac_6 + \sqrt{a^2c_6^2 + ac_6\alpha(t)x_4[c_6 + \alpha(t)x_4(1 - q_{64}(t))]} }{c_6 + \alpha(t)x_4(1 - q_{64}(t))} \quad (19)$$

The specific expression for $x_6^*(t)$ of Eq. (19) does not add new insight, but it is worth noting that once $x_6(t)$ exceeds $x_6^*(t)$, it is possible for $q_{66}(t)$ to start decreasing again without risking to increase translation traffic beyond a . As a matter of fact, it is even possible, although obviously not necessarily practical or desirable, to lower IPv6 quality back below that of IPv4 (but not below that of translation devices). This is because once the IPv6 user base is large enough, this alone is sufficient to entice most ICPs to become IPv6 accessible even if connectivity quality is lower than with IPv4 (but better than what would be realized by going through translation devices). An illustration of this possibility is provided in Fig. 3.

The main findings from the above discussion are summarized in Proposition 2.

Proposition 2 *ISPs can control the volume of IPv6 traffic that undergoes translation (to IPv4) by adjusting the relative connectivity quality of IPv6 and translation compared to that of IPv4. Some of the trade-offs this involves are as follows*

- *If an ISP is unable to offer IPv6 connectivity quality that is better than what it offers through translation, translation traffic volume will keep increasing to eventually exceed any provisioned translation capacity, *i.e.*, $\lim_{x_6 \rightarrow \infty} T(t) = \infty$ if $q_{66}(t) \leq q_{64}(t)$.*
- *If translation quality is low enough, it is possible to keep the volume of translation traffic below a certain threshold without ever making IPv6 connectivity quality better than that of IPv4.*
- *Conversely, if translation quality is too close to that of the current IPv4 connectivity and translation capacity is limited, it is eventually necessary for IPv6 connectivity quality to exceed that of IPv4 to keep the volume of translation traffic below capacity.*
- *Once IPv6 connectivity quality exceeds that of IPv4, increasing the number of IPv4 users that are IPv6 enabled has a positive impact on keeping translation traffic volume bounded. In particular, once a sufficient number of IPv4 users are IPv6 accessible, both upper bounding translation traffic and eventually ensuring complete migration to IPv6 can be realized with IPv6 quality equal but no better than IPv4 quality.*
- *Conversely, once IPv6 penetration has reached a certain level (either by making IPv4 users IPv6 accessible or through growth of the population of IPv6 users), it is possible to reduce IPv6 connectivity quality below that of IPv4 (but above that of translation) while still keeping translation traffic bounded.*

3.2 Model 2 – Connectivity Quality has Precedence

In this section, we briefly consider Model 2, where IPv4 users with IPv6 addresses are selective enough to enable IPv6 connectivity only when it offers higher quality, *i.e.*, $q_{66}(t) > 1$. Because the main results of Section 3.1.2 are mostly insensitive⁸ to the differences between Models 1 and 2, we only explore how IPv6 adoption (by ICPs) is affected by the assumptions of Model 2.

⁸This is because we assume the ISP would realize the negative effect of providing existing IPv4 users with IPv6 connectivity when $q_{66}(t) < 1$, and would therefore set $\alpha(t) = 0$.

3.2.1 ICP IPv6 Adoption

The main difference with the scenario of Model 1 is the absence of IPv6 adoption disincentives associated with a positive $\alpha(t)$, when IPv6 connectivity is worse than that of IPv4. From Eqs. (4), (6), and (7), together with the assumption that the θ_i are uniformly distributed in $[0, 1]$, we can again characterize $\beta(t)$, the level of IPv6 adoption by ICPs,

$$\beta(t) = \begin{cases} 0 & \text{if } q_{66}(t) < q_{64}(t) < 1 & \text{(a)} \\ \left(\frac{x_6(t)(q_{66}(t) - q_{64}(t))}{c_6} \right)^{[0,1]} & \text{if } q_{64}(t) \leq q_{66}(t) < 1 & \text{(b)} \\ \left(\frac{x_6(t)(q_{66}(t) - q_{64}(t)) + \alpha(t)x_4(q_{66}(t) - 1)}{c_6} \right)^{[0,1]} & \text{otherwise} & \text{(c)} \end{cases} \quad (20)$$

Eqs. (20) and (9) differ only in Eq. (20)(b) that accounts for the fact that IPv4 users with IPv6 addresses will not select IPv6 connectivity unless it offers better quality. As a result, there is no negative impact of the ISP's choice of $\alpha(t) > 0$ on ICP IPv6 adoption decisions in this configuration. All the findings of Proposition 1 remain therefore valid, except for those related to providing existing IPv4 users with an IPv6 address when IPv6 quality is worse than that of IPv4. This is summarized in the mostly self-evident next proposition.

Proposition 3 *When users configured with both IPv4 and IPv6 addresses are able to select the connectivity option with the highest quality, increasing the number of such users has no negative impact on the decision by ICPs to become IPv6 accessible. However, for such an increase to have a positive effect on ICPs IPv6 adoption, the quality of IPv6 connectivity must be higher than that of the current IPv4 Internet.*

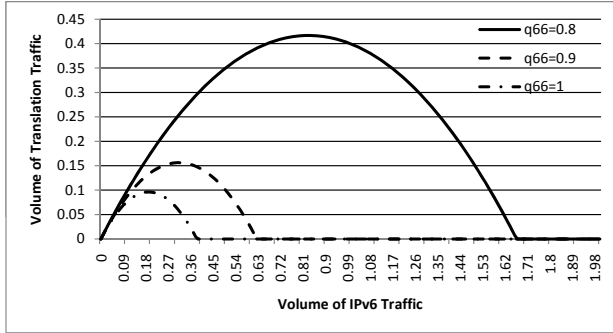


Figure 1: Impact of IPv6 Quality on Translation Traffic Volume ($q_{64} = 0.74$).

3.3 Numerical Examples

This section provides a few representative illustrations of the paper's findings. It assumes $x_4 = 1$, *i.e.*, the number and traffic of IPv4 users is normalized to 1, and $a = 0.1$ and $c_6 = 0.1$. In other words, translation devices have been provisioned to handle a traffic volume equal to 10% of the current IP4 traffic, and the cost of IPv6 provisioning is in the worst case ($\theta_i = 1$) equal to 10% of an ICP's revenue. Given these values, we consider two possible translation quality values, $q_{64} = 0.74$ and 0.79 , where the first value satisfies Eq. (12), namely, $q_{64} = 0.74 < 1 - \frac{c_6}{4a} = 0.75$, but the latter doesn't. The impact of this difference is illustrated by com-

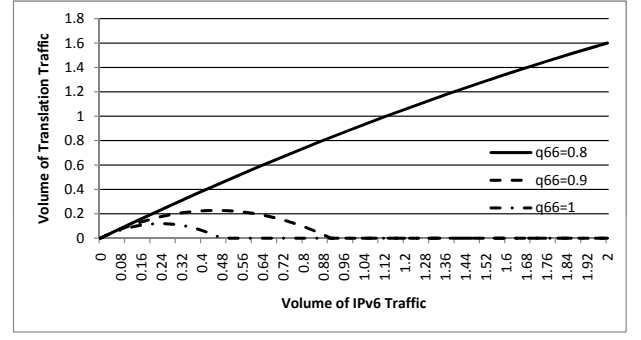


Figure 2: Impact of IPv6 Quality on Translation Traffic Volume ($q_{64} = 0.79$).

paring Figs. 1 and 2, with the former exhibiting consistently lower volumes of translation traffic. Additionally, Fig. 1 also shows that as indicated by Eq. (11), it is possible to keep $T(t) \leq a = 0.1$ with an IPv6 quality no better than that of IPv4 (Eq. (11) states that $q_{66}^{\min} = q_{64} + 1 - \frac{c_6}{4a} = 0.99$ and the figure plots translation traffic for $q_{66} = 1$). Figs. 1 and 2 also illustrate that improving translation quality calls for improving IPv6 quality if one is to keep translation traffic below the provisioned capacity. A similar conclusion applies to ensuring full migration of the IPv4 Internet to IPv6, *e.g.*, when $q_{64} = 0.74$, β reaches 1 for $x_6 = 1.67$ and $q_{66} = 0.8$, but increasing q_{64} to 0.79 calls for correspondingly increasing q_{66} to 0.85 to achieve the same result.

Another perspective on the impact of q_{64} on q_{66} is illustrated in Fig. 3 that plots as a function of the number of IPv6 users, the minimum required IPv6 quality to keep $T(t) \leq a = 0.1$. The figure also shows that once the number of IPv6 users exceeds the value of Eq. (19), it is possible to decrease IPv6 quality, even below that of IPv4, without risking exceeding the capacity of translation devices.

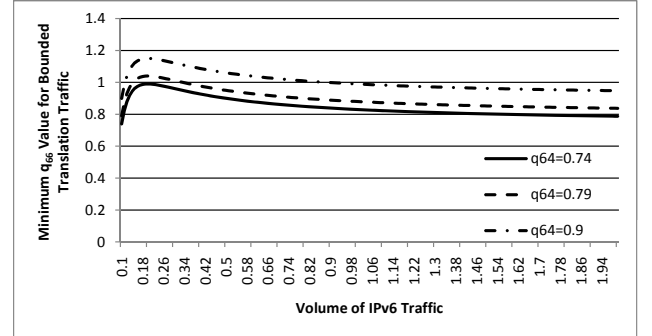


Figure 3: Minimum IPv6 Quality to Keep Translation Traffic Volume Below Provisioned Capacity.

4. CONCLUSION

This paper used a simple model to explore how quality and capacity of translation devices and IPv6 quality could affect both migration of the IPv4 Internet to IPv6, and the volume of traffic that translation devices need to handle.

In spite of its simplicity and obvious limitations, the model helped elucidate a number of interesting issues, and in par-

ticular the often ambiguous role of translation devices. Those devices are mandatory to let IPv6 users access the current Internet that is accessible mostly only over IPv4, and their quality needs to be high enough to satisfy those users. On the other hand, if their quality is too high, it will likely slow down an eventual migration to IPv6. Another more intuitive finding is the role of IPv6 connectivity quality. In particular, if IPv6 quality remains below what is achievable through translation, migration to IPv6 will not occur and the volume of translation traffic will keep growing. Additionally, if translation quality is high enough, *e.g.*, as may be required to entice users to initially accept IPv6, it may be necessary to improve IPv6 quality beyond that of the current IPv4 Internet, if one is to keep translation traffic volume below the provisioned capacity of translation devices. Conversely, increasing the capacity of those devices can facilitate migration to IPv6 without requiring that IPv6 quality exceeds that of IPv4.

The model also established the potentially significant benefits that ISPs can derive from making IPv6 connectivity quality *better* than that of the current IPv4 Internet, at least initially when IPv6 adoption by ICPs is low. A high-quality IPv6 Internet, and obviously publicizing this higher quality, will entice ICPs to make themselves IPv6 accessible and keep translation traffic low. The resulting savings will likely offset costs incurred in improving IPv6 quality. In other words, the higher the initial quality of IPv6, the earlier ICPs are likely to become IPv6 accessible⁹ without waiting for $x_6(t)$ to be large. This will in turn ensure that the volume of translation traffic remains small.

This being said, there are clearly challenges in translating this into reality, the least of which is the nature of the Internet that is made-up of a collection of interconnected autonomous entities. The failure of a few of them to offer high-quality IPv6 connectivity could affect end-to-end IPv6 quality for many users and ICPs. On the other hand, if the larger ISPs, which to some extent stand to gain the most from a faster migration to IPv6, lead the way in offering high(er) quality IPv6 connectivity, this may be sufficient incentive (the corresponding user base is large) for many ICPs to consider making themselves IPv6 accessible.

There are obviously many extensions one could consider to make the simple model used in this paper more realistic. For example, one aspect that the current model ignores is how competition among ICPs and possibly ISPs might affect IPv6 adoption decisions. Other extensions of interest include allowing heterogeneity in content popularity and revenues as well as in connectivity quality. Clearly, the conversion to IPv6 of a highly popular web site has a larger impact on translation traffic volume than that of a less popular site. This will introduce differences, although we don't expect them to drastically affect the findings of the simple model, *e.g.*, the adoption of IPv6 by a popular site is akin similar adoption decisions by several less popular sites. Similarly, the notion of a single quality value for all users and ICPs is clearly an over-simplification, as users experience different levels of connectivity quality when accessing different ICPs, and conversely different users often see different connectivity quality when accessing the same ICP. Exploring models that allow such heterogeneity may be of interest, even if as with ICP heterogeneity we expect the main find-

⁹This obviously reflects the growth of ICP revenues with user activity, which itself grows with network quality.

ings of our simple model to remain mostly valid, *e.g.*, based on an "average" measure of quality.

Another extension of interest, and one that we plan to actively pursue is to minimize the model's reliance on exogenous parameters, and in particular include users among the decision makers, *i.e.*, endogenize $x_6(t)$ by making it a function of the connectivity quality that new users experience. Additionally, it is of interest to capture more accurately the tension between translation costs and connectivity costs, *i.e.*, account for the fact that improving IPv6 connectivity quality is likely to come at a cost. Our goal is to formulate a joint optimization that will incorporate both cost components and seek to characterize the strategy that yields the lowest overall cost. Finally, validating the paper's findings is another worthwhile endeavor, and one of the motivations behind the previously mentioned measurement efforts whose partial results are available at <http://mnlab-ipv6.seas.upenn.edu/monitor/index.html>.

Acknowledgments

The authors would like to acknowledge A. Durand's inputs on characteristics and limitations of translation devices, which helped crystallize the paper's focus. They would also like to thank B. Edelman for insightful comments on an early version of the paper. Finally, they would like to acknowledge the role of the WIE'09 Workshop¹⁰ in initiating discussions that in part motivated the paper.

5. REFERENCES

- [1] Alexa Top Web Sites. Available at <http://www.alexa.com/topsites>.
- [2] K. Das. US Government using IPv6. Available at <http://ipv6.com/articles/general/US.Government.IPv6.htm>.
- [3] S. Deering and R. Hinden. Internet Protocol, version 6 (IPv6) specification. RFC 1883, December 1995.
- [4] R. Draves. Default address selection for Internet Protocol version 6 (IPv6). RFC 3484, February 2003.
- [5] A. Durand, Ed. Dual-stack lite broadband deployments post IPv4 exhaustion. Internet Draft (draft-ietf-softwire-dual-stack-lite-02), October 2009. (Work in progress).
- [6] B. Edelman. Running out of numbers: Scarcity of IP addresses and what to do about it. In *Proc. First Conference on Auctions, Market Mechanisms and Their Applications*. Springer-Verlag's Lecture Notes of the Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering, 2009.
- [7] G. Huston. Is the transition to IPv6 a "market failure"? Available at <http://www.potaroo.net/ispcol/2009-09/v6trans.html>, September 2009.
- [8] C. Labovitz. The end is near, but is IPv6? Available at <http://asert.arbornetworks.com/2008/08/the-end-is-near-but-is-ipv6/>, August 2008.
- [9] Technical and economic assessment of Internet Protocol version 6 (IPv6). US Department of Commerce - Available at <http://www.ntia.doc.gov/ntiahome/ntiageneral/ipv6/final/ipv6final.pdf>, January 2006.

¹⁰See <http://www.caida.org/workshops/wie/0909/>.